AFRL-AFOSR-VA-TR-2015-0398

(YIP) Detecting, Analyzing, Modeling Adversarial Propaganda in Social Media

James Caverlee
TEXAS ENGINEERING EXPERIMENT STATION COLLEGE STATION

10/26/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
14-10-2015

2. REPORT TYPE
Final Performance Report

3. DATES COVERED *(From - To)*
July 2012 - July 2015

4. TITLE AND SUBTITLE
(YIP) Detecting, Analyzing, Modeling Adversarial Propaganda in Social Media

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-12-1-0363

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
James Caverlee
Associate Professor
Department of Computer Science and Engineering
Texas A&M University

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Texas Engineering Experiment Station
1470 William D Fitch Pkwy
College Station, TX 77845-4645

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 North Randolph Street
Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A - Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This research project has focused on detecting, analyzing, and modeling campaigns engaged in strategic manipulation in social media. Research questions investigated include: Can we detect campaigns engaging in strategic manipulation efficiently and effectively, considering the massive and growing scale of social systems? What approaches are in use? How do they evolve? What impacts do spatial and temporal constraints of social media have on strategic manipulation? What are the dynamics of strategic manipulation?

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** *(Include area code)*

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATE COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33315-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report. e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

AFOSR Final Performance Report

Project Title:	(YIP) Detecting, Analyzing, Modeling Adversarial Propaganda in Social Media
Award Number:	FA9550-12-1-0363
Start Date:	July 2012
Program Manager:	Benjamin A. Knott, PhD Air Force Office of Scientific Research Program Officer - Trust and Influence 875 N. Randolph St. Arlington, VA 22203 Phone: 703-696-1142 email: benjamin.knott.2@us.af.mil
Principal Investigator:	Prof. James Caverlee Department of Computer Science and Engineering Texas A&M University College Station, TX 77843-3112 email: caverlee@cse.tamu.edu

Summary

This project has investigated strategic manipulation and adversarial propaganda in social media through a multi-faceted campaign-oriented framework. A campaign is a coherent and organized effort to spread a particular message – e.g., misinformation, propaganda, payloads containing malware, phishing URLs – and stands in contrast to "organic" (or grassroots) efforts. This project has viewed campaigns from multiple perspectives, including content-based (e.g., bag-of-words, LIWC, or other markup), behavioral indicators (e.g., posting patterns, clicking patterns), user demographics (e.g., gender, age), link-based (e.g., friendship network, retweet network), and spatial-temporal footprint in several venues, including Twitter, Fiverr, and Amazon. The overall research goals of this project have been threefold:

- **Detection**: Can we detect campaigns engaging in strategic manipulation efficiently and effectively considering the massive and growing scale of social systems?
- **Analysis**: What approaches are in use? How do they evolve? What impacts do spatial and temporal constraints of social media have on strategic manipulation?
- **Modeling**: What are the dynamics of strategic manipulation? E.g., can we model "tipping points" for campaigns? Does competition impact the reach of a campaign?

A critical challenge for the investigation of strategic manipulation is the lack of ground truth – that is, the knowledge of whether a collection of social media posts is part of a coherent and organized campaign is necessarily hidden from us. Since the origins of strategic manipulation are hidden from us, this project has made critical strides in identifying the ground truth of known strategically manipulated campaigns through new approaches and in exploiting these advances for new analysis and modeling:

- **Bottom-Up Campaign Detection and Analysis.** Targeting strategic crowdsourced manipulation (ala Chinese Online Water Army) where the command-and-control of strategically-organized campaigns can be monitored, models developed for identifying evidence of campaigns, and countermeasures potentially deployed for disrupting the reach and effectiveness of campaigns.
- **Top-Down Campaign Detection and Analysis.** By mining billions of tweets for evidence of strategic manipulation through multiple methods (including content-based and behavior-based techniques) so that candidate campaigns can be extracted from the massive scale of the social web, and subsequently models may be developed for distingishing between organic and strategically organized campaigns.
- **Discovery and Tracking of Topic-Sensitive Opinion Bias.** Given that we may be able to detect strategic manipulation, to what degree does it actually impact opinions? Hence, we have developed methods that – given a topic of interest (e.g., ebola, Syria) and just a handful of seeds to characterize the topic space – semi-automatically discover and track the bias themes associated with opposing sides of a topic, identify strong partisans who drive the online discussion, and infer the opinion bias of "regular" social media participants.
- **Data-Driven Models.** Finally, to simulate large-scale social systems based on parameters derived from real social systems, explore what-if scenarios, and test hypotheses of campaign success and reach, we have created data-driven models based on real system traces. For example, with a comprehensive evaluation over a Twitter system trace, we have evaluated the effectiveness of campaign detectors deployed based on the first moments of a bursting phenomenon in a real system, providing a shield for unsuspecting social media users.

Findings

This project has investigated strategic manipulation and adversarial propaganda in social media through a multi-faceted campaign-oriented framework. A campaign is a coherent and organized effort to spread a particular message – e.g., misinformation, propaganda, payloads containing malware, phishing URLs – and stands in contrast to "organic" (or grassroots) efforts. This project has viewed campaigns from multiple perspectives, including content-based (e.g., bag-of-words, LIWC, or other markup), behavioral indicators (e.g., posting patterns, clicking patterns), user demographics (e.g., gender, age), link-based (e.g., friendship network, retweet network), and spatial-temporal footprint in several venues, including Twitter, Fiverr, and Amazon. In the following, we highlight advances made in this project. References at the conclusion contain detailed information supporting this report.

1. Bottom-Up Campaign Detection and Analysis

The origins of strategic manipulation are necessarily hidden from us – that is, if a message or campaign is the outgrowth of a hidden agenda, how can this hidden agenda be known? Without such access to the origins of strategic manipulation, algorithmic approaches to detecting them must be built over external labels (e.g., a handful of human raters can grade a message as being manipulated or not) which may not be generalizable, may miss out on many manipulated campaigns, and may lead to errors in the assessments of campaigns.

In this project, we have made the first effort toward identifying known manipulated campaigns through the opportunistic mining and analysis of crowdsourcing platforms that often serve as the launching point for strategically manipulated campaigns. The crowdsourcing movement has spawned a host of successful efforts that organize large numbers of globally-distributed participants to tackle a range of tasks, including crisis mapping (e.g., Ushahidi), translation (e.g., Duolingo), and protein folding (e.g., Foldit). Alongside these specialized systems, we have seen the rise of general-purpose crowdsourcing marketplaces like Amazon Mechanical Turk and Crowdflower that aim to connect task requesters with task workers, toward creating new crowdsourcing systems that can intelligently organize large numbers of people. However, these positive opportunities have a sinister counterpart. In particular, we highlight the challenge of *weaponized crowdsourcing*, in which malicious requesters misuse this openness to post tasks that spread malicious URLs in social media, form artificial grassroots campaigns, spread rumor and misinformation, and manipulate search engines. In the same vein, unethical workers will perform these tasks, often by propagating manipulated contents to target sites such as social media sites, search engines, and review sites, resulting in the degradation of information quality and the integrity of these online communities.

To illustrate, Figure 1 shows a typical workflow, wherein (1) a requester first posts one of these tasks (here, a "crowdturfing" task), (2) identifies the appropriate workers to complete this task, and (3) finally, these workers spread their misinformation in a target venue like a social network, a forum/review site, a search engine, or blog. An example of a crowdturfing task description that we sampled from the crowdsourcing platform Microworkers.com requires workers to have at least 50 Twitter followers, search for a certain keyword on Google, and then click on a website in the search results. In addition, it requires the workers to retweet an article in the website to Twitter. This task targets not only a search engine but also a social media site, hoping to boost the target website's rank by artificially manipulating both a search engine and a social network. At the time of our collection, 222 workers had completed this task for $0.60 per task completion.

Twitter. In our first effort, we focus on uncovering the ecosystem of strategically manipulated campaigns in the Twitter social media environment.

- We first analyze the types of malicious tasks and the properties of requesters and workers on Western crowdsourcing sites such as Microworkers.com, ShortTask.com and Rapidworkers.com. Previous

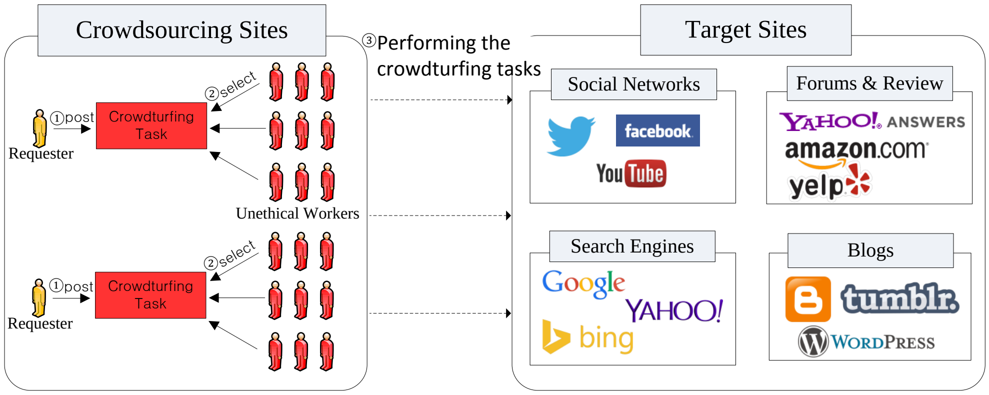

Figure 1: The interactions between malicious requesters and unethical workers.

researchers have investigated Chinese-based crowdsourcing sites; to our knowledge this is the first study to focus primarily on Western crowdsourcing sites.

- Second, we propose a framework for linking tasks (and their workers) on crowdsourcing sites to social media, by monitoring the activities of social media participants on Twitter. In this way, we can track the activities of crowdturfers in social media where their behavior, social network topology, and other cues may leak information about the underlying crowdturfing ecosystem.

- Based on this framework, we identify the hidden information propagation structure connecting these workers in social media, which can reveal the implicit power structure of crowdturfers identified on crowdsourcing sites. Specifically, we identify three classes of crowdturfers – professional workers, casual workers, and middlemen – and we demonstrate how their roles and behaviors are different in social media.

- Finally, we propose and develop statistical user models to automatically differentiate among regular social media users and workers. Our experimental results show that these models can effectively detect previously unknown Twitter-based workers.

To illustrate this final point, as time passes, a pre-built classifier can lose its classification accuracy because crowdturfing workers may change their behavioral patterns to hide their true identities from the classifier. Hence, we tested a classifier built over a previous point to determine if it is still effective at a later point in time. We created Twitter campaigns a month later in three crowdsourcing sites – Microworkers.com, ShortTask.com and Rapidworkers.com – to collect new workers' Twitter account information consisting of their profile information, tweets and following-follower information. As shown in Table 1, we collected 368 Twitter user profiles and their recent 200 messages (in total, 40,344 messages).

Table 1: A New Worker Dataset.

Class	User Profiles	Tweets
Workers	368	40,344

Next, we evaluated our previously built classifier, with this dataset as the testing set, by measuring how many workers in the set are correctly predicted. Table 2 presents its experimental result. It confirms that our classifier is still effective even with the passage of time with 94.3% accuracy, 0.971 F_1 measure and 0.057 false negative rates.

Table 2: Worker detection over the new dataset using the original detection model.

Classifier	Accuracy	F_1	FNR
Random Forest	94.3%	0.971	0.057

In summary, this positive experimental result shows that our classification approach is promising to find new workers in the future. Our proposed framework linking crowdsourcing workers to social media works effectively. Even though workers may change memes or URLs which they want to spread as the time passes, their behaviors and observable features such as activity patterns and linguistic characters will be similar and are different from regular users.

Fiverr. We next turned to Fiverr – a fast growing micro-task marketplace and the 125th most popular site (Alexa 2013) – to be the first to answer the following questions: what are the most important characteristics of buyers (a.k.a. customers) and sellers (a.k.a. workers)? What types of tasks, including crowdturfing tasks, are available? What sites do crowdturfers target? How much do they earn? Based on this analysis and the corresponding observations, can we automatically detect these crowdturfing tasks? Can we measure the impact of these crowdturfing tasks? Can current security systems in targeted sites adequately detect crowdsourced manipulation?

To answer these questions, we made the following contributions:

- First, we collect a large number of active tasks (these are called gigs in Fiverr) from all categories in Fiverr. Then, we analyze the properties of buyers and sellers as well as the types of crowdturfing tasks found in this marketplace. To our knowledge, this is the first study to focus primarily on Fiverr.
- Second, we conduct a statistical analysis of the properties of crowdturfing and legitimate tasks, and we build a machine learning based crowdturfing task classifier to actively filter out these existing and new malicious tasks, preventing propagation of crowdsourced manipulation to other web sites. To our knowledge this is the first study to detect crowdturfing tasks automatically.
- Third, we feature case studies of three specific types of crowdturfing tasks: social media targeting gigs, search engine targeting gigs and user traffic targeting gigs.
- Finally, we purchase active crowdturfing tasks targeting a popular social media site, Twitter, and measure the impact of these tasks to the targeted site. We then test how many crowdsourced manipulations Twitter's security can detect, and confirm the necessity of our proposed crowdturfing detection approach.

Amazon. Finally, we investigate how crowdsourcing platforms have created a new attack vector for polluting online reviews. These platforms enable a large-scale, potentially difficult-to-detect workforce of deceptive review writers. To illustrate, we have found an example task that asks each worker to leave a high product rating and a "skeptical and realistic but positive" review. Compared to traditional spam bots that typically leave identifiable footprints, these human-powered deceptive reviews are inherently distinct, linked only by their common theme and not in common keywords, phrases, or other easily identifiable signals. And since crowdsourced deceptive reviews are generated by humans rather than bots, their ongoing detection is even more challenging since crowds can actively circumvent detection methods.

Unfortunately, there is a significant gap in our understanding of crowdsourced manipulation of online reviews and effective methods for uncovering such manipulation, due to a number of challenges. One critical challenge is a lack of clear ground truth for analyzing deceptive reviews and in building countermeasures; it is difficult to ascertain which reviews are deceptive and which ones are legitimate. A second important

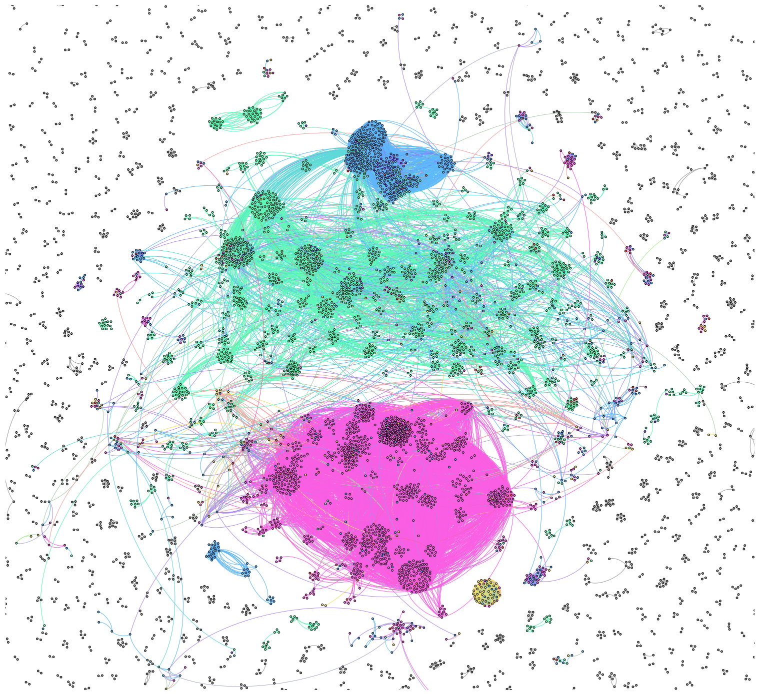

Figure 2: Detected clusters in the Amazon reviewer-reviewer graph. Nodes and edges are colored based on the cluster to which they were assigned. Two out of the three largest clusters are deceptive reviewer clusters.

challenge is that polluters may rely on multiple communication channels to coordinate their activities – including private methods such as email and instant messenger which are difficult to observe – and so deceptive intent may be further obscured. To tackle these issues, we proposed a three-part effort:

- First, we propose a novel sampling method for identifying products that have been targeted for manipulation and a seed set of deceptive reviewers who have been enlisted through the command-and-control of crowdsourcing platforms. To our knowledge, this is the first effort toward "pulling back the curtain" to uncover clear evidence linking deceptive intent with actual reviews.

- Second, we augment the seed set of sampled deceptive reviewers to identify additional deceptive reviewers who participate as part of hidden groups of coordinated manipulation. To capture the hidden infrastructure underlying deceptive reviews, we exploit connections between reviewers through a reviewer-reviewer graph clustering approach based on a conditional random field that models individual potentials (of single reviewers) in combination with pair potentials (between two reviewers).

- Finally, we embed the results of this probabilistic model into a classification framework for detecting crowd manipulated reviews. We find that the proposed approach achieves up to 0.96 AUC, outperforming both traditional detection methods and an alternative SimRank-based clustering approach. An example cluster of deceptive reviewers is highlighted in Figure 2.

Our classification approach using the reviewer clustering results as a feature significantly outperformed a classification approach not using the reviewer clustering results. Specifically, our approach with clustering results have achieved 0.77 AUC in unbalanced testing set and 0.96 AUC in balanced testing set, improving 54% AUC in unbalanced testing set and 8% AUC in balanced testing set respectively compared with the classification approach without using reviewer clustering results. Additionally, the proposed approach has outperformed a SimRank and K-medoids based approach in terms of effectiveness and efficiency.

2. Top-Down Campaign Detection and Analysis.

In contrast to the bottom-up approach – in which crowd workers are first identified for known ground truth of strategic manipulation – we additionally have investigated top-down approaches for mining billions of social media posts to uncover evidence of campaigns. These efforts have been along two primary dimensions – content-based methods for identifying linkages across "free text" posts over billions of messages and in behavioral-based methods for identifying linkages across the posting patterns of users.

Content-Based. First, we focus on detecting one particular kind of coordinated campaign – those that rely on "free text" posts, like those found on blogs, comments, forum postings, and short status updates (like on Twitter and Facebook). For our purposes, a campaign is a collection of users and their posts bound together by some common objective, e.g., promoting a product, criticizing a politician, or inserting disinformation into an online discussion. Our goal is to link messages with common "talking points" and then extract multi-message campaigns from large-scale social media. Detecting these campaigns is especially challenging considering the size of popular social media sites like Facebook and Twitter with 100s of millions of unique users and the inherent lack of context in short posts.

- Concretely, we propose and evaluate a content-based approach for identifying campaigns from the massive scale of real-time social systems. The content-driven framework is designed to effectively link free text posts with common "talking points" and then extract campaigns from large-scale social media. Note that text posts containing common "talking points" means the contents of the posts are similar or the same.
- We find that over millions of Twitter messages, the proposed framework can identify 100s of coordinated campaigns, ranging in size up to several hundred messages per campaign. The campaigns themselves range from innocuous celebrity support (e.g., fans retweeting a celebrity's messages) to aggressive spam and promotion campaigns (in which handfuls of participants post hundreds of messages with malicious URLs).
- Through an experimental study over millions of Twitter messages we identify five major types of campaigns – Spam, Promotion, Template, News, and Celebrity campaigns – and we show how these campaigns may be extracted with high precision and recall. We also find that the less organic campaigns (e.g., Spam and Promotion) tend to be driven by a higher ratio of messages to participants (corresponding to a handful of accounts "pumping" messages into the system).
- Based on this observation, we propose and evaluate a user-centric campaign detection approach. By aggregating the messages posted by a single user, we find that the method can successfully discover cross-user correlations not captured at the individual message level (e.g., for two users posting a sequence of correlated messages), resulting in more robust campaign detection. In addition, we analyze each campaign type's temporal behavior to see the possibility to automatically determine a campaign's campaign type.

To illustrate, we highlight our work on a *user-aggregated* perspective. In the message level campaign detection, we have viewed all messages without consideration for *who* is posting the messages. By also considering user-level information, we are interested to see how this impacts campaign detection. The intuition is that by aggregating the messages posted by a single user, we may discover cross-user correlations not captured at the individual message level (e.g., for two users posting a sequence of correlated messages), leading to more robust campaign detection.

Definition 1 (User-Aggregated Message Graph) *A user-aggregated message graph is a graph $G_u = (V, E)$ where V is a collection of n users' aggregate messages $V = \{M_{u_1}, M_{u_2}, ...M_{u_n}\}$. An edge $(M_{u_i}, M_{u_j}) \in$*

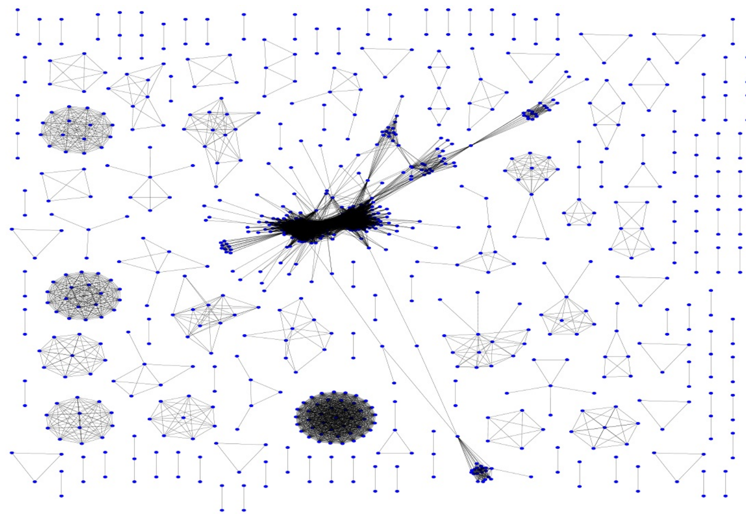

Figure 3: 303 candidate campaigns in the user-aggregated message graph.

E exists for every pair of vertices (M_{u_i}, M_{u_j}) in V where $confidence$ $(Mu_i, Mu_j) > threshold$, for some measure of confidence and threshold. In the confidence computation, message similarity for every pair of messages (m_{ix}, m_{jy}) is computed where $corr(m_{ix}, m_{jy}) > \tau$, $m_{ix} \in M_{u_i}$, $m_{jy} \in M_{u_j}$ and $M_{u_i}, M_{u_j} \subseteq M$, for some measure of correlation and some parameter τ.

An important challenge is to define the correlation across vertices in the user-aggregated message graph, since each vertex now represents multiple messages (and so straightforward adoption of the message-level correlation approach is insufficient). To compute user-based correlation, we propose a measure called *confidence* that aggregates message-message correlation and reflects (i) that one edge in a one-to-many match receives same weight comparing to the edge in a one-to-one edge; (ii) that extra edges in a one-to-many match receive less weight than the weight for the edge in a one-to-one match, but still credits the one-to-many match for more evidence of user-based correlation.

Concretely, we calculate confidence in the following way: Given two users u_1 and u_2 and their latest *k* messages M_{u_i} = $\{m_{i1}$, m_{i2}, ..., $m_{ik}\}$ where *i* is a user id (i.e., 1 or 2 in our example). First, we compute pairwise message correlation across M_{u_1} and M_{u_2}, where pairs are $P = \{m_{1x}, m_{2y} | 1 \leq x, y \leq k\}$. If the correlation of a pair in *P* is larger than threshold τ, we consider the pair to be correlated. By continuing this procedure for each pair in *P*, we have correlated pairs P' and can calculate: (1) the number of pairs in P', $N = |\{m_{1x}, m_{2y} | corr(m_{1x}, m_{2y}) \geq \tau, 1 \leq x, y \leq k\}|$; and (2) the minimum n between number of distinct messages belonging to P' in M_{u_1} and number of distinct messages belonging to P' in M_{u_2}, where $n = MIN(|\{m_{1x}|\ m_{1x} \in M_{u_1}$ and $M_{1x} \in P'\}|,|\{m_{2y}|\ m_{2y} \in M_{u_2}$ and $m_2 j \in P'\}|)$. Now, we define that *confidence* as:

$$confidence = \alpha n + (1 - \alpha)(N - n)$$

where α is the weight for the only edge in a one-to-one match or one edge in a one-to-many match, and $1 - \alpha$ is the weight for each of the extra edges in a one-to-many match. We assigned 0.95 to α to balance between αn and $(1 - \alpha)(N - n)$. By applying this method, we uncover the 303 campaigns as shown in Figure 3.

Behavior-Based. To complement the content-based approach, we additionally investigate the potential of *behavioral analysis* for uncovering which URLs are campaign-oriented and which are not. By behavioral signals, we are interested both in the aggregate behavior of *who is posting* these URLs in social systems

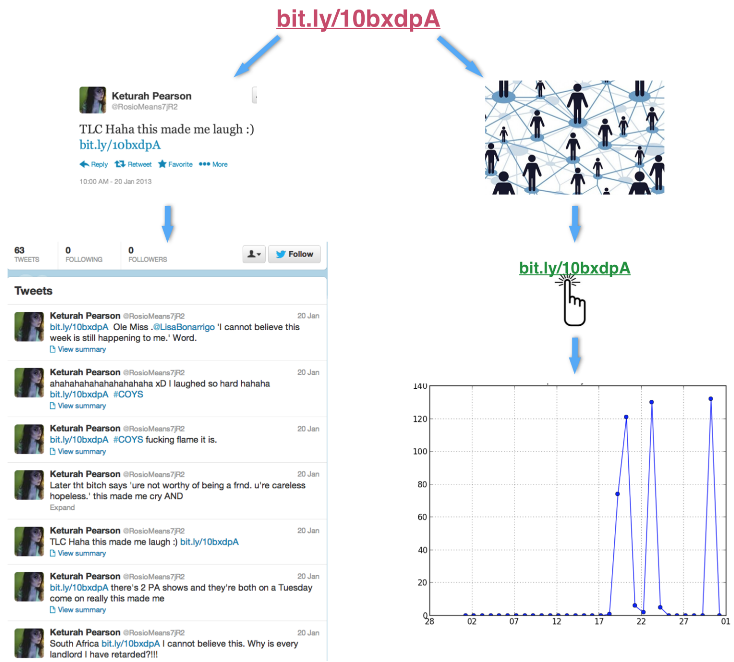

Figure 4: Studying spam URL detection in social media from two perspectives: (i) Posting behavior (left); (ii) Click behavior (right)

and *who is clicking* on these URLs once they have been posted. These behavioral signals offer the potential of rich contextual evidence about each URL that goes beyond traditional detection methods that rely on blacklists, the content of the URL, its in-links, or other link-related metadata. Unfortunately, it has historically been difficult to investigate behavioral patterns of posts and clicks. First, many social systems provide restricted (or even no) access to posts, like Facebook. Second, even for those systems that do provide research access to a sample of its posts (like Twitter), it has been difficult to assess how these links are actually received by the users of the system via clicks. As a result, much insight into behavioral patterns of URL sharing has been limited to proprietary and non-repeatable studies.

- Hence, we have begun a behavioral examination of campaign-oriented URL detection through two distinct perspectives (see Figure 4): (i) the first is via a study of how these links are posted through publicly-accessible Twitter data; (ii) the second is via a study of how these links are received by measuring their click patterns through the publicly-accessible Bitly click API.

- Concretely, we propose and evaluate fifteen click and posting-based behavioral features, including: for postings – how often the link is posted, the frequency dispersion of when the link is posted (e.g., is it posted only on a single day in a burst? or is it diffusely posted over a long period?), and the social network of the posters themselves; and for clicks – we model the click dynamics of each URL (e.g., does it rapidly rise in popularity? are there multiple spikes in popularity?) and consider several click-related statistics about each URL, including the total number of clicks accumulated and the average clicks per day that a URL was actually clicked.

- Through extensive experimental study over a dataset of 7 million Bitly-shortened URLs posted to Twitter, we find that these behavioral signals provide overlapping but fundamentally different perspectives on URLs. Through this purely behavioral approach for spam URL detection, we can achieve high precision (0.86), recall (0.86), and area-under-the-curve (0.92). Compared to many existing methods

that focus on either the content of social media posts or the destination page – which may be easily manipulated by spammers to evade detection – this behavior-based approach suggests the potential of leveraging these newly-available behavioral cues for robust, on-going spam detection.

3. Discovery and Tracking of Topic-Sensitive Opinion Bias

Given that we may be able to detect strategic manipulation, to what degree does it actually impact opinions? Social media has increasingly become a popular and important platform for "regular" people to express their opinions, without the need to rely on expensive and fundamentally limited conduits like newsstandss and broadcast television. These opinions can be expressed on a variety of themes including politically-charged topics like abortion and gun control as well as fun (but heated) rivalries like android vs. iOS and Cowboys vs. 49ers. Our interest is in creating a flexible tool for discovering and tracking the themes of opinion bias around these topics, the strong partisans who drive the online discussion, and the degree of opinion bias of "regular" social media participants, to determine to what degree particular participants support or oppose a topic of interest.

However, assessing topic-sensitive opinion bias is challenging. First, the opinion bias of "regular" users may not be as pronounced as prominent figures, so discerning this bias will require special care. Second, how opinion bias manifests will inevitably vary by topic, so a system should be adaptable to each topic. Third, the themes by which people express their opinions may change over time depending on the circumstances (e.g., gun control debates may take different forms based on the ebb and flow of elections, recent shooting incidents, and so forth). As a result, assessing bias should be adaptive to these temporal changes.

Hence, we develop a lightweight system – BiasWatch – for discovering and tracking opinion bias in social media. BiasWatch begins by taking just two hand-picked seeds to characterize the topic-space (e.g., "pro-choice" and "pro-life" for abortion) as *weak labels* to bootstrap the opinion bias framework. Concretely, we leverage these hand-picked seeds to identify other emerging (and often unknown) themes in social media, reflecting changes in discourse as new arguments and issues arise and fade from public view (e.g., an upcoming election, a contentious news story). We propose and evaluate two approaches for expanding the hand-picked seeds in the context of Twitter to identify supporting and opposing hashtags – one based on co-occurrence and one on signed information gain. We use these discovered hashtags to identify strong topic-based partisans (what we dub *anchors*). Based on the social and information networks around these anchors, we propose an efficient opinion-bias propagation method to determine user's opinion bias – based on both content and retweeting similarity – and embed this method in an optimization framework for estimating the topic-sensitive bias of social media participants.

- First, we build a systematic framework – BiasWatch – to discover biased themes and estimate user-based opinion bias quantitatively under the context of controversial topics in social media. We propose an efficient optimization scheme – called User-guided Opinion Propagation [UOP] – to propagate opinion bias. By feeding just two opposing hashtags, the system can discover bias-related hashtags, find bias anchors, and assess the degree of bias for "regular" users who tweet about controversial topics.

- Second, we evaluate the estimation of users' opinion bias by comparing the quality of the proposed opinion bias approach versus several alternative approaches over multiple Twitter datasets. Overall, we see a significant improvement of 20.0% in accuracy and 28.6% in AUC on average over the next-best method.

- Third, we study the effect of different approaches for biased theme discovery to measure the impact of newly discovered biased hashtags as additional supervision. We observe that the newly discovered

Table 3: Comparison of performance with alternative opinion bias estimators. Boldface: the best result for each topic among all methods. '$*$' marks statistically significant difference against the best of alternative opinion bias estimators (with two sample t-test for $p \leq 0.05$).

Method	Accuracy				AUC			
	gun control	abortion	obamacare	average	gun control	abortion	obamacare	average
SWN	0.560	0.527	0.465	0.517	0.570	0.531	0.541	0.547
uCC	0.534	0.537	0.516	0.529	0.533	0.527	0.522	0.527
uCCL	0.586	0.530	0.520	0.545	0.584	0.531	0.546	0.554
wSVM+IS	0.696	0.825	0.786	0.769	0.745	0.790	0.594	0.710
wSVM+SIG	0.860	0.884	0.727	0.824	0.844	0.874	0.800	0.839
UOP*	0.851	0.847	0.826	0.841	0.853	0.843	0.842	0.846
LCGC+SIG	0.858	0.900	0.811	0.856	0.857	0.900	0.864	0.874
UOP†	0.881	0.906	0.894	0.894	0.861	0.903	0.915	0.893
UOP	**0.908***	**0.915**	**0.945***	**0.923***	**0.883***	**0.910**	**0.945***	**0.913***

hashtags are often associated with the underlying community of similar opinion bias, and that they temporally fluctuate due to the impact of new controversial events.

- Finally, we demonstrate how these inferred opinion bias scores can be integrated into user recommendation by giving similar-minded users a higher ranking. We show that the integration can improve the recommendation performance by 26.3% in precision@20 and 13.8% in MAP@20. This result implicitly confirms the principle of homophily in the context of opinion bias, and demonstrates how topic-sensitive opinion bias can enrich user modeling in social media.

We compare our proposed BiasWatch framework with the alternative opinion bias estimators. For uCCL, the weight-balancing parameter θ is selected from $\{0.05, 0.1, 0.15, 0.2, 0.25\}$ to be 0.1 for best performance. Other parameter settings in UOP*, UOP† and UOP are the same as in previous experiments. We choose the best seed expansion approach (via SIG) for all methods. The final results are shown in Table 3.

Overall, user-guided approaches give much better performance than unsupervised methods, indicating that by just a small amount of human guidance — two opposite seed hashtags, the performance can be boosted significantly. Moreover, UOP gives the best performance, reaching an average accuracy and AUC of 0.923 and 0.913, respectively (an improvement of 20.0% and 28.6% over supervised baseline wSVM+IS). Note that the sentiment based approach SWN gives unsatisfactory results, probably because user's opinion bias is multifaceted and can be reflected by the topical arguments or factual information published by the user. For example, one of the anti-obamacare tweets reads, "Double Down: Obamacare Will Increase Avg. Individual-Market Insurance Premiums By 99% For Men, 62% For Women. #Forbes". Also, we can see that UOP† gives better results than LCGC+SIG, indicating that our framework works better in capturing user's opinion bias. UOP, however, gives better performance than UOP†, confirming that initial bias anchors determined through biased hashtags are noisy, and that UOP is able to correct some wrongly determined bias anchors due to the l-1 norm regularization on ideal bias scores.

Furthermore, we see from the table that UOP† has an improvement of 0.053 and 0.047 for accuracy and AUC over UOP*, respectively. Compared to the corresponding improvement of 0.016 and 0.027 by uCCL over uCC, it is considerably higher. This indicates that retweeting links are more effective in contributing to bias propagation with the help of bias anchors. When some users are correctly "labeled" by discovered biased hashtags, opinion bias can be propagated more effectively through retweeting links.

4. Data-Driven Models

Finally, we have created models and simulations of large-scale social systems based on parameters derived from real social systems, so that we can explore what-if scenarios, and test hypotheses of campaign success and reach. Testing hypotheses without such models and simulations is challenging. Large-scale social systems are typically proprietary and responsible to their current user base, so it is infeasible to automatically "stress-test" such a system by subjecting it to hundreds or thousands of malicious users. An alternative is to take a representative snapshot of a system and measure the current level of threats in the system and characterize their reach and effectiveness. However, this approach alone may not be suitable for understanding the system's future state, as social systems are constantly evolving. Hence, we take a two fold approach. First, we take a *data-driven modeling* approach, in which we simulate a large-scale social system based on parameters derived from a real system. In this way, we can vary system parameters – like the fraction of malicious users in the system, their strategies, and the countermeasures available to system operators – to explore the resilience of these systems. We pair the data-driven model with a *comprehensive evaluation over a Twitter system trace*, in which we evaluate the effectiveness of countermeasures deployed based on the first moments of a bursting phenomenon in a real system.

We consider a social system of interest $\mathcal{S}$, consisting of a set of content items $\mathcal{C}$ (e.g., videos, tweets, etc.), a set of topics $\mathcal{T}$ for which each content item is associated (e.g., the "London Olympics" topic, the "Steve Jobs" topic, etc.), and a set of users $\mathcal{U}$, who participate in the system by posting and viewing content items. For example, a user in $\mathcal{U}$ may post a tweet "Thank you #SteveJobs The world will miss you", where the tweet is associated with the topic indicated by the hashtag #SteveJobs. Similarly, a user may post a video to YouTube associated with the "London Olympics" topic by including a tag or descriptive text at the time of upload. We use the symbols u, c, and t to denote a user in $\mathcal{U}$, a content item in $\mathcal{C}$, and a topic in $\mathcal{T}$.

To populate a social system, we initialize the system with a set of topics and a set of users. To model users in a social system, we define two sets of users: a good user set U^+ and a bad user set U^-. Good users post content items that are associated with a "correct" topic. Bad users, on the other hand, post content items that are irrelevant to the topic they are associated with. For example, a bad user may post a spam video, but intentionally describe it as a "London Olympics" video. When users post to the system, we assume they have access to both the set of topics $\mathcal{T}$ as well as the current subset of "popular" topics $\mathcal{T}_{pop}$ (in practice, these popular topics may be known to users via prior knowledge or explicitly advertised by the system, as in the case of Twitter trending topics or popular YouTube videos). The system proceeds in step-wise fashion; at each time increment, users generate content items according to a particular posting model. Good users post according to the *good user model*:

Good User Model:

for each user $u \in U^+$ do
  with probability γ^+ decide to post:
    with probability δ^+:
      select a popular topic $t \in \mathcal{T}_{pop}$ and relevant item c;
    else:
      select at random a topic $t \in \mathcal{T}$ and relevant item c.

At each time increment, a good user chooses to post something with the user content generation probability γ^+. If a user decides to post a content item, an already popular topic is selected with probability δ^+; alternatively, the user decides to post to a random topic. A bad user follows a similar process, but always posts spam content items:

Bad User Model:

for each user $u \in U^-$ do

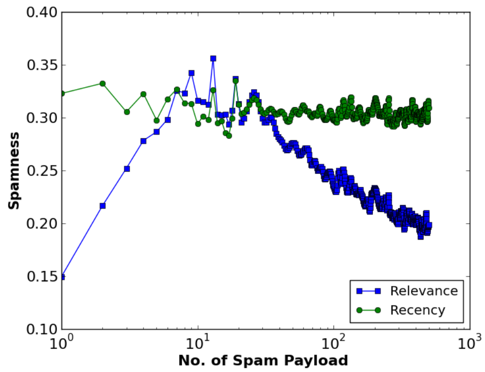

Figure 5: Coordinated Spam: By focusing their efforts, groups can achieve even higher impact.

with probability γ^- decide to post:
 with probability δ^-:
  a popular topic $t \in \mathcal{T}_{pop}$ and spam item c;
 else:
  select at random a topic $t \in \mathcal{T}$ and spam item c.

Notice that the user content generation probability γ and the popular topic probability δ may vary between the good and bad user models. As part of our data-driven simulation, we will vary these parameters to reflect different spammer behaviors. For example, a spammer may adopt a high rate of content generation relative to good users (e.g., $\gamma^- \gg \gamma^+$) in an attempt to flood the system with spam content. Alternatively, a spammer seeking to maximize their potential audience may choose to focus only on popular topics and so adopt a popular topic probability much greater than the good user model (e.g., $\delta^- \gg \delta^+$).

Given the approach for populating a social system, we now consider how users access the content posted in the system. We assume that users monitor topics by one of two methods:

- *By recency*: In the first access model, users interested in a topic access the k-most recently posted items related to the topic. This recency approach is akin to the "Most Recent Uploads" functionality on YouTube, viewing comments associated with a blog by their posting order (from recent to oldest), and Twitter's basic search.

- *By relevance*: The second access model imposes a relevance ordering over content items associated with a topic. This relevance-based approach may incorporate the popularity of an item (e.g., rank images in order of the number of clicks they have accumulated), content and link-based ranking (e.g, applying IR principles), or learning-to-rank methods. For modeling purposes, we assume that content items are ranked by their occurrence count, with all duplicates removed to maintain diversity (i.e., item c_i posted 20 times is ranked first; item c_j posted 10 times is ranked second; and so on).

By instantiating this model, we can consider different what-if scenarios and evaluate various countermeasures. For example, assume we have multiple participants coordinating to spread a particular campaign item (e.g., a spam message). We can consider a coordinated spam approach in which spammers are assigned to a group which is associated with a common pool of spam payloads. For the following experiment, we assume that spammers share a common pool of spam payloads, and we vary the number of spam payloads.

Using this coordinated approach, we observe in Figure 5 that the recency-based approach is largely unaffected, but that it remains highly susceptible to spam. The relevance-based approach shows that spammers have a potential "sweet spot" for targeting spam. At a low number of payloads, the spamness is relatively low since the spammers promote a few payloads which possibly pollute one or two out of the top-k results.

As the number of payloads increases, the coordinating spam group can achieve an impact equal to or even better than under the recency-based approach. However, as the number of payloads continues to increase, the effectiveness for the coordinating spam group falls, because the power promoting payloads is distributed across too many payloads, meaning no single one can penetrate the top-k, and hence be exposed to end users interested in the topic.

Publications

Peer-Reviewed Archival Publications

- K. Lee, P. Tamilarasan, and J. Caverlee. Crowdturfers, Campaigns, and Social Media: Tracking and Revealing Crowdsourced Manipulation of Social Media. Seventh International AAAI Conference on Weblogs and Social Media (ICWSM), 2013. [acceptance rate = 20%]
- K. Lee, K. Y. Kamath, and J. Caverlee. Combating Threats to Collective Attention in Social Media: An Evaluation. Seventh International AAAI Conference on Weblogs and Social Media (ICWSM), 2013. [acceptance rate = 20%]
- K. Lee, J. Caverlee, Z. Cheng, and D. Z. Sui. Campaign Extraction from Social Media. ACM Transactions on Intelligent Systems and Technology (TIST), 2013.
- K. Lee, S. Webb, and H. Ge. The Dark Side of Micro-Task Marketplaces: Characterizing Fiverr and Automatically Detecting Crowdturfing. Eighth International Conference on Weblogs and Social Media (ICWSM), 2014. [acceptance rate = 23%]
- C. Cao and J. Caverlee. Detecting Spam URLs in Social Media via Behavioral Analysis. 37th European Conference on Information Retrieval (ECIR), 2015. [acceptance rate = 23%]
- A. Fayazi, K. Lee, J. Caverlee, and A. Squicciarini. Uncovering Crowdsourced Manipulation of Online Reviews. 38th Annual ACM SIGIR Conference, 2015. [acceptance rate = 20%]
- C. Cao, J. Caverlee, K. Lee, H. Ge, and J. Chung. Organic or Organized? Exploring URL Sharing Behavior. 24th ACM International Conference on Information and Knowledge Management (CIKM), 2015. [acceptance rate = 18%]
- H. Lu and J. Caverlee. BiasWatch: A Lightweight System for Discovering and Tracking Topic-Sensitive Opinion Bias in Social Media. 24th ACM International Conference on Information and Knowledge Management (CIKM), 2015. [acceptance rate = 18%]

Workshop Papers and Short Papers

- C. Cao and J. Caverlee. Behavioral Detection of Spam URL Sharing: Posting Patterns versus Click Patterns. IEEE/ACM International Conference on Advances in Social Network Analysis and Mining (ASONAM), 2014.

Book Chapters

- J. Caverlee and K. Lee. Weaponized Crowdsourcing: An Emerging Threat and Potential Countermeasures. In Transparency in Social Media (Springer) edited by S. Matei, M. Russell, and E. Bertino. 2014.

Outreach and Impact

Graduates:

Kyumin Lee (PhD), now Assistant Professor at Utah State University
Amir Fayazi (MS), now Software Engineer at Google
Plus multiple PhD and MS students still in-progress.

- Tutorial on Social Spam, Campaigns, Misinformation, and Crowdturfing. 23rd International World Wide Web Conference (WWW), 2014.
- Tutorial on Social Media Threats and Countermeasures. 8th International Conference on Weblogs and Social Media (ICWSM), 2014.
- Panelist at CollaborateCom 2013 in Austin, October 21, 2013.
- Invited speaker at the NSF kredible.net Reputation, Trust and Authority Workshop at Stanford University on October 18, 2013.

1.

1. Report Type

Final Report

Primary Contact E-mail

Contact email if there is a problem with the report.

caverlee@cse.tamu.edu

Primary Contact Phone Number

Contact phone number if there is a problem with the report

979-209-9998

Organization / Institution name

Texas A&M University

Grant/Contract Title

The full title of the funded effort.

(YIP) Detecting, Analyzing, Modeling Adversarial Propaganda in Social Media

Grant/Contract Number

AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-12-1-0363

Principal Investigator Name

The full name of the principal investigator on the grant or contract.

James Buchanan Caverlee

Program Manager

The AFOSR Program Manager currently assigned to the award

Benjamin A. Knott

Reporting Period Start Date

07/15/2012

Reporting Period End Date

07/14/2015

Abstract

This project has investigated strategic manipulation and adversarial propaganda in social media through a multi-faceted campaign-oriented framework. A campaign is a coherent and organized effort to spread a particular message -- e.g., misinformation, propaganda, payloads containing malware, phishing URLs -- and stands in contrast to ``organic'' (or grassroots) efforts. This project has viewed campaigns from multiple perspectives, including content-based (e.g., bag-of-words, LIWC, or other markup), behavioral indicators (e.g., posting patterns, clicking patterns), user demographics (e.g., gender, age), link-based (e.g., friendship network, retweet network), and spatial-temporal footprint in several venues, including Twitter, Fiverr, and Amazon. The overall research goals of this project have been threefold:

Detection: Can we detect campaigns engaging in strategic manipulation efficiently and effectively considering the massive and growing scale of social systems?

Analysis: What approaches are in use? How do they evolve? What impacts do spatial and temporal constraints of social media have on strategic manipulation?

Modeling: What are the dynamics of strategic manipulation? E.g., can we model ``tipping points'' for

campaigns? Does competition impact the reach of a campaign?

A critical challenge for the investigation of strategic manipulation is the lack of ground truth -- that is, the knowledge of whether a collection of social media posts is part of a coherent and organized campaign is necessarily hidden from us. Since the origins of strategic manipulation are hidden from us, this project has made critical strides in identifying the ground truth of known strategically manipulated campaigns through new approaches and in exploiting these advances for new analysis and modeling:

Bottom-Up Campaign Detection and Analysis. Targeting strategic crowdsourced manipulation (ala Chinese Online Water Army) where the command-and-control of strategically-organized campaigns can be monitored, models developed for identifying evidence of campaigns, and countermeasures potentially deployed for disrupting the reach and effectiveness of campaigns.

Top-Down Campaign Detection and Analysis. By mining billions of tweets for evidence of strategic manipulation through multiple methods (including content-based and behavior-based techniques) so that candidate campaigns can be extracted from the massive scale of the social web, and subsequently models may be developed for distinguishing between organic and strategically organized campaigns.

Discovery and Tracking of Topic-Sensitive Opinion Bias. Given that we may be able to detect strategic manipulation, to what degree does it actually impact opinions? Hence, we have developed methods that -- given a topic of interest (e.g., ebola, Syria) and just a handful of seeds to characterize the topic space -- semi-automatically discover and track the bias themes associated with opposing sides of a topic, identify strong partisans who drive the online discussion, and infer the opinion bias of ``regular'' social media participants.

Data-Driven Models. Finally, to simulate large-scale social systems based on parameters derived from real social systems, explore what-if scenarios, and test hypotheses of campaign success and reach, we have created data-driven models based on real system traces. For example, with a comprehensive evaluation over a Twitter system trace, we have evaluated the effectiveness of campaign detectors deployed based on the first moments of a bursting phenomenon in a real system, providing a shield for unsuspecting social media users.

Distribution Statement

This is block 12 on the SF298 form.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your SF298 form. A blank SF298 can be found here. Please do not password protect or secure the PDF The maximum file size for an SF298 is 50MB.

SF 298.pdf

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

final-report.pdf

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

Peer-Reviewed Archival Publications

K. Lee, P. Tamilarasan, and J. Caverlee. Crowdturfers, Campaigns, and Social Media: Tracking and Revealing Crowdsourced Manipulation of Social Media. Seventh International AAAI Conference on Weblogs and Social Media (ICWSM), 2013. [acceptance rate = 20\%]

K. Lee, K. Y. Kamath, and J. Caverlee. Combating Threats to Collective Attention in Social Media: An Evaluation. Seventh International AAAI Conference on Weblogs and Social Media (ICWSM), 2013. [acceptance rate = 20\%]

K. Lee, J. Caverlee, Z. Cheng, and D. Z. Sui. Campaign Extraction from Social Media. ACM Transactions on Intelligent Systems and Technology (TIST), 2013.

K. Lee, S. Webb, and H. Ge. The Dark Side of Micro-Task Marketplaces: Characterizing Fiverr and Automatically Detecting Crowdturfing. Eighth International Conference on Weblogs and Social Media (ICWSM), 2014. [acceptance rate = 23\%]

C. Cao and J. Caverlee. Detecting Spam URLs in Social Media via Behavioral Analysis. 37th European Conference on Information Retrieval (ECIR), 2015. [acceptance rate = 23\%]

A. Fayazi, K. Lee, J. Caverlee, and A. Squicciarini. Uncovering Crowdsourced Manipulation of Online Reviews. 38th Annual ACM SIGIR Conference, 2015. [acceptance rate = 20\%]

C. Cao, J. Caverlee, K. Lee, H. Ge, and J. Chung. Organic or Organized? Exploring URL Sharing Behavior. 24th ACM International Conference on Information and Knowledge Management (CIKM), 2015. [acceptance rate = 18\%]

H. Lu and J. Caverlee. BiasWatch: A Lightweight System for Discovering and Tracking Topic-Sensitive Opinion Bias in Social Media. 24th ACM International Conference on Information and Knowledge Management (CIKM), 2015. [acceptance rate = 18\%]

Workshop Papers and Short Papers

C. Cao and J. Caverlee. Behavioral Detection of Spam URL Sharing: Posting Patterns versus Click Patterns. IEEE/ACM International Conference on Advances in Social Network Analysis and Mining (ASONAM), 2014.

Book Chapters

J. Caverlee and K. Lee. Weaponized Crowdsourcing: An Emerging Threat and Potential Countermeasures. In Transparency in Social Media (Springer) edited by S. Matei, M. Russell, and E. Bertino. 2014.

Changes in research objectives (if any):

Change in AFOSR Program Manager, if any:

Extensions granted or milestones slipped, if any:

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, $K)

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Oct 13, 2015 09:10:14 Success: Email Sent to: caverlee@cse.tamu.edu